Workshop on Material Appearance Modeling (2018)
H. Rushmeier and R. Klein (Editors)

Towards Practical Rendering of Fiber-Level Cloth Appearance Models

Adrian Alejandre Carlos Aliaga[†] Julio Marco Adrian Jarabo Adolfo Muñoz

Universidad de Zaragoza, I3A

Abstract

Accurate representation of realistic cloth appearance is of high importance in many industry fields such as entertainment and textile design. However, microstructure of fibers and their optical properties generate very complex lighting effects, often not reproduced by empirical and theoretical models. In contrast, data-driven appearance models obtained with simulations on explicit representations of fibers have proved to yield accurate cloth appearance, but resulting in discretized distribution functions that require costly precomputations, massive storage, and expensive evaluation in render time. Finding efficient representations for these models is therefore of key importance to find good trade-offs between accuracy and affordable computational costs. In this work we explore the use of different analytical models to represent these data-driven distributions, which arises as a promising middle-ground solution to this problem with benefits in both storage, computational cost, and affordable generation of new fiber appearance models. We base this analysis on our recent work where we provide highly detailed tabulations of different types of cloth fibers appearance. We analyze the spectral component of different fiber appearance functions, and observe that just ten levels of spherical harmonics are sufficient to represent the appearance many smooth fibers. We also propose a generic method to fit Gaussian Mixture Models to massively tabulated appearance functions, reducing storage costs from hundreds of MBs to a few KBs, and producing equivalent results 40 times faster. We finally analyze how interpolations in the space of fibers absorption can be exploited to generate novel fiber appearance functions without requiring costly brute-force precomputations.

1. Introduction

Cloth is ubiquitous in real world, specially in man-made scenarios. As such, notable effort has been made to render it in an accurate manner. However, modeling the volumetric appearance of cloth is very challenging, since it depends on both macro- and microscopic levels of appearance, ranging from the type of knitting pattern, the type of yarn being used and, at the smallest scale, the appearance of the fibers making such yarns. Detailed volumetric representations have been proposed to accurately represent the macroscopic knit [ZJMB11,KSZ*15], and the structure of yarns [ZLB16]. These highly-detailed representations allow for high-quality depiction of cloth. Unfortunately, these works still relied on simple *bidirectional curve scattering distribution functions (BCSDFs)*, that had no connection with the low-level optical properties of cloth fibers.

Recently, we presented an appearance model for cloth fibers [ACG*17], that accounted for the singular appearance of individual fibers. We based on accurate measurements of the structure of fibers, based on SEM images of the fiber's cross-section, and optical properties obtained from measurements of the fibers roughness and the molecular absorption of fibers dyes. Such low-level description of fibers allowed highly detailed BCSDFs for cloth fibers, including complex anisotropic reflection not captured by parametric models, while establishing a direct relationship between low-level manufacture parameters of cloth fibers, potentially helping as a tool for predictive rendering of cloth.

Unfortunately, this accuracy comes at the price of expensive precomputation and massive storage due to the need of high-resolution 4D tabulated BCSDFs. In order to build the BCSDFs, costly Monte Carlo simulations are run to mimic a virtual goniometer of the fiber's reflectance field. While this is done as a pre-process, explicit light transport simulation is required for each cross section and *depth of shade* (amount of dye on the fiber). Moreover, this simulation results into a dense, high-resolution tabulated representation of the BCSDF, that needs to be accessed in run-time requiring a large memory footprint, and that makes sampling an expensive task. These two problems reduce the applicability of Aliaga et al.'s BCSDFs in practical renders.

In this work, we report additional analyses on the BCSDFs developed by Aliaga et al., and describe our efforts to partially solve the two main problems on such highly detailed scattering functions. We first analyze the spectral content of precomputed cloth BCSDFs by using their spherical harmonics (SH) expansion: We found that up to 30 levels of SH (900 coefficients) are required to accurately represent all details of the very anisotropic cloth like silk, but just

---

[†] Carlos Aliaga is now with Desilico Labs.

DOI: 10.2312/mam.20181195

ten levels are required for accurately depicting smoother fibers like cotton.

Then, we propose a generic method to compactly represent tabulated radiance functions without sacrificing accuracy by using *Gaussian Mixture Models* (GMM). Our goal is to use GMMs to obtain a compact and analytic version of fiber BCSDFs that can be used in render time, avoiding prohibitive storage costs and providing an easy way to perform importance sampling for efficient evaluation during Monte Carlo integration.

Finally, we describe our initial explorations on how to generate novel BCSDFs from Aliaga's precomputed ones. We analyze the space defined by the depth of shade (DoS, i.e. fibers absorption), and find that we can create new BCSDFs by simple interpolation of three precomputed ones, by carefully sampling the space of DoS.

2. Textile Fibers BCSDFs

Several different analytical scattering models have been proposed to model the appearance of fibers [MJC*03, ZW07, dFH*11, KM16]. However, they rely on the assumption that fibers have simple shape. Moreover, they do not provide any connection between fabrication parameters and the optical properties of fibers, so the appearance of fibers needs to be specified manually or via optimization. Aliaga et al. [ACG*17] developed a low-level model based on precise measurements of real-world fibers structure and composition. Based on such measurements, the authors created virtual replicas of the fibers, and computed their 4D scattering field via Monte Carlo simulation. This resulted into highly detailed tabulated BCSDFs, with a resolution of 2 degrees on each dimension.

These functions provide accurate results (see Figure 1, left), at the expense of high storage costs, usually taking more than 500 MB of storage per fiber function for a single depth of shade. Moreover, for efficient integration into a general render engine they required a tabulation-based sampling procedure, requiring tabulating a precomputed 2D CDF per incoming direction, and a costly search-based importance sampling. In the following, we analyze Aliaga's BCSDFs, and propose a novel representation based on GMMs that allow compact storage, and efficient importance sampling.

3. Compact BCSFs using Gaussian Mixture Models

In our work we propose a generic method to create a compact representation of tabulated radiance functions without sacrificing accuracy by using *Gaussian Mixture Models* (GMM). A GMM is a probabilistic model based on the linear combinations of k Gaussians. Our goal is to use GMMs to obtain a compact and analytic version of textile fibers BCSDFs that can be used in render time, avoiding prohibitive storage costs and providing an easy way to perform importance sampling for Monte Carlo integration.

Estimating the GMM Among several algorithms to train a GMM, the *Expectation-Maximization* (EM) algorithm refines the parameters of a GMM to increase the log-likelihood of the model to the provided input samples. In the original algorithm, the distribution function is estimated by drawing samples from this function, and feeding the sample coordinates to the EM algorithm for iterative refining. To estimate 4D tabulated functions with GMMs we rely on a modified version of the EM algorithm [Kaw08]. Instead of feeding the algorithm with a set of samples proportional to the radiance intensity of the tabulated function, we feed the coordinate of each cell of the table only once, and use the radiance value at the cell as a weight to indicate the density of the function at that point. In addition, we can interpret our GMM as a probability density function by normalizing the GMM weights, ensuring they add up to 1. This serves to efficiently generate random samples following the distribution of the original tabulated function, for example to perform multiple importance sampling.

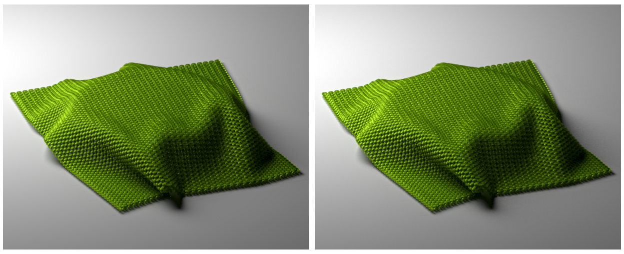

Figure 1: *Render comparison of cotton fabric. The left image was rendered using the original tabulated cotton BCSDF of Aliaga et al. [ACG*17] and uniform sampling, with 16000 samples per pixel during 59 hours. The right image use our GMM representation of the BCSDF and GMM-based importance sampling, with 512 samples per pixel, taking only 1.4 hours.*

Results As shown in an equal-quality comparison in Figure 1, by approximating the tabulated BCSDF of cotton fabric [ACG*17] with our GMMs we are able to faithfully render a cotton patch by reducing render time over 40 times (from 59 hours to 1.4 hours) thanks to efficient importance sampling. Additionally, our model provides a dramatic reduction on the storage requirements of the measured appearance functions. Representing the tabulated BCSDF with a set of 2D GMMs reduces the size from the original 526 MB tabulated function to just 140 KB. As we can see in Figure 1, generating samples with our Gaussian Mixture Models of the cotton fabric we have render a fabric in 1.8 hours instead of the 59 hours required to obtain an image with similar quality using uniform sampling.

The representation quality of this approach depends mainly of two factors. The first one, the number of 2-dimensional Gaussian Mixture Models in which we are dividing the tabulated function, since a low number of GMMs will fail to capture high frequencies. The second factor is the approximation accuracy due to using regular Gaussians to approximate a BCSDF parameterized in spherical coordinates. This shows some singularities at the edges of the 2D slices, where theta and phi have their limits value (see Figure 2). An interesting solution for this problem would be to use a different kind of probabilistic functions but defined in a spherical domain (e.g. spherical Gaussians [WRG*09, IDN12]), or using alternative parameterizations of the BCSDF.

Using GMMs to represent appearance functions has interesting properties which can be used in future research. Since performing convolutions between Gaussians has closed form, it could serve to

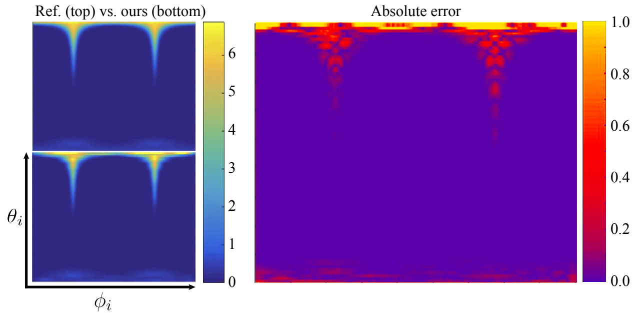

Figure 2: *Comparison between a 2D slice of a polyester BCSDF and the Gaussian Mixture Model approximation with 50 Gaussians. The first column shows the original tabulated function (top) and their GMM approximation (bottom). Due to the structure of Gaussians the error is concentrated at the superior edge of the image, where* θ_i *is equal to* π*. The right image shows the absolute error of the approximation.*

combine multiple appearance functions represented by GMMs, or account for the directional distribution function of fibers in voxel-based volumetric models. This is a promising option to improve the sampling of complex materials.

4. Analyzing the Space of DoS

A key problem of the approach of Aliaga et al. is the need to precalculate via simulation for every single fiber cross section and depth of shade (DoS). This is not only very expensive in terms of precomputation, but also requires storing a different BCSDF for each color, making very expensive the use of e.g. textures. To alleviate this issue, we analyze the effect of the DoS on the appearance of BCSDFs. In particular, we evaluate the error introduced by obtaining a BCSDF f_ξ for a DoS ξ by linearly interpolating two BCSDFs f_0 and f_1 with DoS given by ξ_0 and ξ_1. We compute the BCSDF for a range of ten DoSs, for set of different fibers. Then, we compute the error introduced by interpolating between three anchor points. Figure 3 shows the result of this analysis. What we observe is that using just three anchor points ξ_0, ξ_i and ξ_1 is enough for low error in the full range. In addition, we found that the intermediate anchor point ξ_i needs to be close to the smallest value of depth of shade (i.e. less absorbing fiber, Figure 3 right).

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (CHAMELEON project, grant agreement No 682080), DARPA (project REVEAL), and the Gobierno de Aragón.

References

[ACG*17] ALIAGA C., CASTILLO C., GUTIERREZ D., OTADUY M. A., LOPEZ-MORENO J., JARABO A.: An appearance model for textile fibers. *Computer Graphics Forum 36*, 4 (2017), 35–45. 1, 2

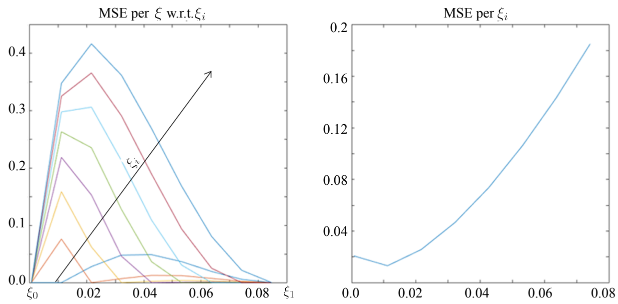

Figure 3: *Error analysis of interpolating fiber BCSDFs with respect to the DoS* ξ*. Left, normalized MSE for each value* ξ *(x-axis) as a function of the anchor point* ξ_i*. Right, average MSE for each anchor point* ξ_i*. We can observe that the error introduced with just three anchor points with* ξ_i *close to* ξ_0 *provides a reasonably good approximation.*

[dFH*11] D'EON E., FRANCOIS G., HILL M., LETTERI J., AUBRY J.-M.: An energy-conserving hair reflectance model. *Computer Graphics Forum 30*, 4 (2011). 2

[IDN12] IWASAKI K., DOBASHI Y., NISHITA T.: Interactive bi-scale editing of highly glossy materials. *ACM Trans. Graph. 31*, 6 (Nov. 2012), 144:1–144:7. 2

[Kaw08] KAWABATA T.: Multiple subunit fitting into a low-resolution density map of a macromolecular complex using a gaussian mixture model. *Biophysical journal 95*, 10 (2008), 4643–4658. 2

[KM16] KHUNGURN P., MARSCHNER S.: Azimuthal scattering from elliptical hair fibers. *ACM Trans. Graph.* (2016). 2

[KSZ*15] KHUNGURN P., SCHROEDER D., ZHAO S., BALA K., MARSCHNER S.: Matching real fabrics with micro-appearance models. *ACM Trans. Graph. 35*, 1 (2015). 1

[MJC*03] MARSCHNER S. R., JENSEN H. W., CAMMARANO M., WORLEY S., HANRAHAN P.: Light scattering from human hair fibers. *ACM Trans. Graph. 22*, 3 (2003). 2

[WRG*09] WANG J., REN P., GONG M., SNYDER J., GUO B.: All-frequency rendering of dynamic, spatially-varying reflectance. In *ACM SIGGRAPH Asia 2009 Papers* (New York, NY, USA, 2009), SIGGRAPH Asia '09, ACM, pp. 133:1–133:10. 2

[ZJMB11] ZHAO S., JAKOB W., MARSCHNER S., BALA K.: Building volumetric appearance models of fabric using micro ct imaging. *ACM Trans. Graph. 30*, 4 (2011). 1

[ZLB16] ZHAO S., LUAN F., BALA K.: Fitting procedural yarn models for realistic cloth rendering. *ACM Trans. Graph. 35*, 4 (2016). 1

[ZW07] ZINKE A., WEBER A.: Light scattering from filaments. *IEEE Trans. Vis. Comp. Graph. 13*, 2 (2007). 2